

PCT

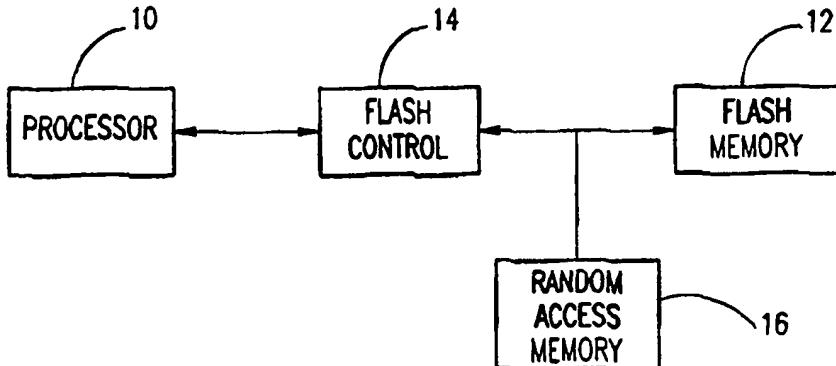
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : G06F 12/02	A1	(11) International Publication Number: WO 94/20906 (43) International Publication Date: 15 September 1994 (15.09.94)
(21) International Application Number: PCT/US94/01848		(81) Designated States: AT, AU, BB, BG, BR, BY, CA, CH, CZ, DE, DK, ES, FI, GB, HU, JP, KP, KR, KZ, LK, LU, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SK, UA, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 28 February 1994 (28.02.94)		
(30) Priority Data: 08/027,131 8 March 1993 (08.03.93) US		
(71) Applicants: M-SYSTEMS LTD. [IL/IL]; P.O.Box 58032, 61580 Tel Aviv (IL). M-SYSTEMS INC. [US/US]; 200 Broadhollow Road, Melville, NY 11747 (US).		Published <i>With international search report.</i>
(72) Inventor: BAN, Amir; 47 Yehuda Hamaccabi, 62309 Tel Aviv (IL).		
(74) Agent: FRIEDMAN, Mark, M.; c/o Sheinbein, Robert, 2940 Birchtree Lane, Silver Spring, MD 20906 (US).		

(54) Title: FLASH FILE SYSTEM



(57) Abstract

The provision of a flash memory (12), virtual mapping system, which includes a flash controller (14) and a random access memory (16) for storing mapping tables, that allows data to be continuously written to unwritten physical address locations. The virtual memory map (14, 16) relates flash memory physical location addresses in order to track the location of data in the memory.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LJ	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

FLASH FILE SYSTEM

BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to an improved system for storing and
5 retrieving information in flash memories, and more particularly to a system
that organizes and manages data written to a flash memory.

Description of the Prior Art

As will be appreciated by those skilled in the art, electrically
erasable and programmable read-only memories (EEPROMs) comprised of
10 flash-type, floating-gate transistors have been described in the art and are
presently commercially available. These so-called flash memories are non-
volatile memories similar in functionality and performance to EPROM
memories, with an additional functionality that allows an in-circuit,
programmable, operation to erase blocks of the memory. In a flash
15 memory, it is not practical to rewrite a previously written area of the
memory without a preceding block erase of the area. While this invention
will be described in the context of a flash memory, those skilled in the art
will understand that its teachings are also applicable to data storage devices
with the same write, read, and block erase before write characteristics as
20 flash memories.

In a typical computer system, the operating system program is
responsible for data management of the data storage devices that are a part
of the system. A necessary, and usually sufficient, attribute of a data
storage device to achieve compatibility with the operating system program
25 is that it can read data from, and write data to, any location in the data
storage medium. Thus, flash memories are not compatible with typical
existing operating system programs, since data cannot be written to an area
of flash memory in which data has previously been written, unless the area
is first erased.

30 Software products have been proposed in the prior art to allow a
flash memory to be managed by existing computer operating programs

without modification of the operating system program. However, these prior art programs operate the flash memory as a "write once read many" device. This prior art software product cannot recycle previously written memory locations. When all locations are eventually written the memory 5 cannot be further used without specific user intervention.

SUMMARY OF THE INVENTION

An object of this invention is the provision of a method (i.e., software, firmware or hardware) to control and manage access to a flash memory so that the flash memory appears to the computer operating 10 system as a data storage device in which it is possible to read data from, and write data to, any flash memory location. A method that allows flash memory to emulate random access memories and allows existing computer operating systems to provide all other required support in the same manner provided by standard random access memories and independent of the 15 emulation method.

Briefly, this invention contemplates the provision of a flash memory, virtual mapping system that allows data to be continuously written to unwritten physical address locations. The virtual memory map relates flash memory physical location addresses in order to track the location of data 20 in the memory.

The flash memory physical locations are organized as an array of bytes. Each of the bytes in the array is assigned a number of address by means of which the byte is physically accessible, referred to herein as the physical address space. Each of the bytes in the array has a second 25 address, called the virtual address space. A table, called a virtual map, converts virtual addresses to physical addresses. Here it should be noted, the virtual address space is not necessarily the same size as the physical address space.

A contiguous, fixed-length group of physical byte addresses from a block. For example, assuming a block size of 512 bytes, a byte with a physical address of 256211 is bytes number 211 in block 500 (256211 : 512 = 500 + 211). One or more physically contiguous flash memory areas 5 (called zones) that can be physically erased using suitable prior art flash memory technology comprise a unit and each unit contains an integral number of blocks.

The virtual memory map is a table in which the first entry belongs to virtual block 0, the second to virtual block 1, and so on. Associated in 10 the table with each virtual block address there is a corresponding physical address. In a read from the flash memory operation, a computer generated address is decoded as a virtual block address and a byte location within the block. The virtual memory map is used to convert the virtual block address to a physical block address; the byte location is the same in the 15 virtual address space and the physical address space.

In a write operation, the computer generated address is again interpreted as a virtual block address and a byte location within the block. The virtual memory map converts this to physical memory block address. If the flash memory block corresponding to the physical address is then 20 currently written, it is generally not possible to write to this physical address. An unwritten block is therefore located and written to. The virtual memory map is changed so that the unwritten physical block address is mapped to the original virtual address and original physical address is denoted as unusable and remains unusable until there is a zone 25 erase operation that erases the unit that includes that block. It will be noted that a write operation assumes that an entire block will be rewritten. This is the manner in which computer systems usually access data in a storage media. However, it will be appreciated that in general, any desired number of bytes could be written to the new storage location.

In a preferred embodiment of the invention, each unit is assigned a logical unit address that remains unchanged as the unit is rewritten into a new physical address location in flash memory. The virtual map contains references to the logical unit addresses rather than the physical unit addresses so that data movement during unit transfers has no effect on the virtual map.

Each unit has a usage map of all the blocks within the unit; the virtual address of a block, if it is mapped, and special characters to denote free blocks and to denote unusable blocks.

10 Unusable blocks of previously written flash memory are reclaimed by transferring memory units that include the unusable blocks to a reserved unwritten space in the flash memory. Only the usable blocks are written in the transfer operation so that, as rewritten, the locations where the unusable blocks were, are not rewritten in the reserved space and are thus 15 usable. After having been rewritten, the original memory unit space is flash erased as a unit and thus becomes an unwritten reserved space to which a subsequent transfer can be made.

Also, in a preferred embodiment of the invention, the virtual map is stored primarily in the flash memory with only a small secondary virtual 20 map in random access memory. The virtual map in flash memory is stored in blocks and organized into pages whose size is equal to the product of the number of bytes in a block times the number of physical block addresses this number of bytes represents. A secondary random access memory contains the page addresses. In reading data for a given virtual 25 address, the page number is determined by dividing address by the page size. The result indexes the secondary virtual map to find the correct primary virtual map block. the remainder is used to calculate the required physical address for the virtual map stored in flash memory. To alter the virtual map in flash memory, the altered map is written into a free block 30 and the secondary map in random access memory is altered to reflect the

change in the primary map location. The replaced block is marked as deleted.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

Figure 1 is a block diagram illustrating functional components of a system in accordance with one embodiment of a system in accordance with the teachings of this invention.

10 Figure 2 is a pictorial illustration of one level of flash memory organization in accordance with the teachings of this invention.

Figure 3 is a pictorial illustration of how a unit is formatted.

Figure 4 is a pictorial representation illustrating how the computer generated addresses are mapped to physical addresses.

15 Figure 5 is a flow diagram illustrating a read operation.

Figure 6 is a flow diagram illustrating a write operation.

Figure 7 is a pictorial diagram illustrating the status of a unit before and after a transfer operation.

Figure 8 is a flow diagram of a transfer operation.

20 Figure 9 is a flow diagram illustrating the operation where a major portion of the virtual to physical map is stored in flash memory.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

Referring now to Figure 1 of the drawings, in a typical system a processor 10, in combination with its operating system software, issues a series of read and write commands to read from, and write data to, specific address locations in a random access memory. As will be appreciated by those skilled in the art, in a random access storage device, such as a disk memory, data can be written to, or read from, any address location. In the

system of Figure 1, the processor 10 writes data to, and reads data from, a flash memory 12 in blocks at specific address locations. Although zones of the flash memory 12 can be erased, currently written address locations cannot be rewritten until the entire zone is erased. In accordance with the 5 teachings of this invention, a flash memory controller 14 provides a fully rewritable virtual address space so that the flash memory 12 emulates a random access memory, such as a disk memory, and the processor operating system software provides all other required operating support (such as a file system) in the same manner as it provides for a standard 10 random access memory, and in a manner that is independent of the flash memory 12 and its controller 14. A typical system also includes a conventional random access memory 16. It will be appreciated that controller 14 functions may be carried out in software, firmware or hardware, and would not necessarily exist as a physically separate unit as 15 the drawing suggests.

Referring now to Figure 2, it illustrates in part the organization of the flash memory. The flash memory has a number of zones labeled here as zone A, zone B, etc. Each zone is comprised of a number of contiguous physical memory locations that can be block erased using conventional, 20 well known, flash memory technology. The zones are organized as units only four of which are shown, labeled in the drawing as; UNIT #1, UNIT #6, UNIT N-1 and TRANSFER UNIT. Each unit is comprised of at least one zone, of a plurality of contiguous zones. Here illustrated, each unit is comprised of two zones (i.e., UNIT #1 - zone A and zone B; UNIT #2 - 25 zone C and zone D, TRANSFER UNIT - zone x2 and 2x).

Each unit is comprised of an integral number of addressable blocks and each block, in turn, is comprised of a contiguous, fixed length group of bytes. At all times, there will be a unit in the memory 12 that is unwritten (i.e., TRANSFER UNIT), so that active blocks in a unit that is 30 to be erased can be written to this unwritten unit prior to erasing the unit.

Referring now to Figure 3, each unit contains an integral number of contiguous data blocks 21 that are in turn comprised of contiguous byte addresses, that can be addressed as a block number and offset within the block. Each block in a unit has a unit can be addressed by block number 5 and offset with the unit. Each unit has a unit header 23 and a map 25 of the allocation status of each block in the unit. The unit header 23 includes a format identifier, and the logical unit number of the unit. Because data must move physically during a unit transfer, the unit number preferably remains unchanged even as the physical location of the unit in the flash 10 memory 12 changes. In addition, the header may also include system-wide information. The block allocation map 25 has a word for each block that denotes its status and its offset in the unit. The status indications are: "block free and writable"; "block deleted and not writable"; "block allocates and contains user data"; and virtual address of the block (back pointer).

15 As previously mentioned, preferably each unit is assigned a logical unit number that does not change, even though the physical location in the memory of the unit change. As illustrated in Figure 4, the computer 10 generated addresses 29 are comprised of a block number and a block offset. These address are interpreted by the flash controller 14 as virtual 20 addresses, and a virtual map is used to establish a correspondence between the virtual address space and physical address space. The virtual map changes as blocks are rewritten an the virtual address space is therefore dynamic. It should be noted also that, at any given time, a block or blocks in the virtual address space may be unmapped to the physical address 25 space, and that blocks in the physical address space may be unwritten and, therefore, free to be written into.

Since the data moves physically during a unit transfer to an unwritten unit space, the units are assigned logical unit numbers that do not change with changes in the physical location of the unit in the memory.

30 A virtual map 31 maps block numbers to logical unit address in the first

step of a two level address translation. The logical unit address is an address relative to a logical unit number, similar to a physical address, which is an address relative to a physical unit number. The logical unit number is the high order binary digits of the logical address and may be 5 derived from the logical address by a bit shift operation. The logical address 33 obtained from map 31 includes a logical unit number along with the offset of the block within the unit.

A logical unit table 35 translates the logical unit number to a physical unit number for the logical unit. This two-step address translation 10 procedure obviates the need to change block addresses in the map when a unit is moved to a new physical location.

In a read operation the virtual address 29 comprised of a block address, for example, initially is mapped to a logical unit number and a block offset within the unit in the addressed block. Map 35 maps the unit 15 number 33 to a physical address 37 for the unit along with the offset of the addressed 37 block within the unit, and the addressed data block is read from this physically location. Here it is assumed that data is read and written on a block basis as is typically done. Of course, data could be written and read on a byte basis using the same principle, if desired. 20 Figure 5 is a flow diagram illustrating this read operation. As previously explained, the virtual address 29 is mapped to a logical address (block 40) in the first step of a two-step address translation. In the second step, the logical address is mapped to a physical address in the flash memory, block 41. Data at this physical address is read, block 42, which terminates this 25 operation.

In a write operation, the virtual address 29 is again mapped initially to a logic unit number and a block offset within the unit. The controller 14 algorithm examines the block allocation map 25 for this unit. If the block corresponding to this address has been written, a write command 30 cannot be executed at the corresponding physical address. The control

algorithm scans the block allocation maps 25 for each unit until a free block is located. The status of the block in the block map 25 at the original unit address is changed to deleted in the block in the allocation map, and the status of the free block is changed to written. The virtual 5 map 31 is updated so that the original virtual address now points to the new logical address where the write operation is to take place. This logical address is mapped to a physical address, in the manner previously described, and the block is written to this address. Figure 6 is a flow diagram illustrating this write operation. In a write operation the virtual 10 address 29 is mapped to a logical unit address, block 45, and the unit allocation for the unit is examined, block 46. If in decision block 47 the unit address is free, the unit address is mapped to a physical address, block 48, and data is written to this physical address, block 49, and the operation ends. If the logical address is not free (block 47), the unit tables are 15 scanned to locate a free address in the unit allocation tables, block 50. This new logical address is mapped to a physical address, block 51, and the data is written to this physical address, block 52. The unit allocation tables are updated (block 53) to indicate that the original block is deleted and not writable, and that the new block is allocated and contains user 20 data. The virtual to logical address map is then updated to point to the new physical address of the data corresponding to the original virtual address, blocks 54 and 55.

To enable reading and writing operations to continue without limitation, physical memory space is reclaimed periodically. As explained 25 previously, at least one unit of the memory is reserved at all times so that it consists entirely of free blocks and serves as a TRANSFER UNIT.

Referring now to figure 7, an active unit is selected (here, UNIT #M) and all of its currently-mapped active blocks are read and then written to the TRANSFER UNIT. The selected unit #M is then block erased, and 30 it becomes the TRANSFER UNIT while the transfer unit to which the

active blocks have been written becomes, in this example, unit #M. Figure 7 illustrates the status of the units before and after a transfer operation. Figure 8 is a flow diagram of this transfer operation. In a transfer operation a unit is selected for transfer, block 60, and the active data 5 blocks in the selected unit are read, block 61. These active data blocks are then written to the transfer unit at positions in the transfer unit corresponding to the positions at which they were located in the original unit, block 62. The original unit selected is then flash erased, block 63, and the logical to physical address map is changed so that the selected unit 10 becomes the transfer unit and the transfer unit is assigned the unit number of the selected unit, block 64.

The system thus far described requires a virtual map whose contents are freely updated, and such a map may be stored in a conventional random access memory. However, assuming, for example, a block size of 15 512 bytes, since the virtual map contains a entry for each block, and each entry may be, for example, 4 bytes long (i.e., capable of addressing up to 4 Gigabytes of memory), a flash memory of 80 Mbytes would require a memory of 640 Kbytes to store the map tables. In order to limit the amount of random access memory required to store the virtual map, in a 20 preferred embodiment of the invention, a major portion of the map data is stored in the flash memory 12 itself, and a secondary virtual map that maps virtual addresses from the computer to the primary virtual map is stored in a random access memory, such as memory 16. An important point to be made here is that the secondary virtual map arrangement makes the 25 procedure for reading and writing the virtual map identical to the procedures for reading and writing ordinary data explained earlier. The virtual map itself treated in a manner equivalent to the user data in the foregoing description and the virtual map stored in random access memory (i.e., secondary virtual map) is the equivalent of the virtual map in the 30 previous description.

In this embodiment, the virtual map resides in the flash memory 12 at negative virtual addresses; ordinary space starts at virtual address zero. The virtual map maps the negative address used by itself, so that the virtual map residing in flash memory can be read and written like ordinary user 5 data, and only the portion of the virtual map that maps itself (i.e., the secondary virtual map) resides in random access memory.

In a simplified example, assume a virtual map of 6000 bytes stored in twelve virtual map blocks, each of 512 bytes. Assuming a four-byte address, each block can store 128 physical addresses. Thus, each block 10 contains the addresses of 64 Kbytes of virtual flash memory. Each block of virtual flash memory addresses is considered as a page and the random access memory stores the page addresses; (in this example, only 48 bytes) mapped to the address blocks. In reading data from a given virtual address, the address is divided by the page size (64 Kbytes) to obtain a 15 page number in the secondary virtual memory that maps to the page block of the primary virtual map where the address is stored. With the virtual memory page block, the procedure to map to a specific flash memory physical address can proceed in the manner already described. For example, after the virtual address is divided by the page size, the remainder 20 can be divided by the virtual memory block size (e.g., 512) to obtain an index to the array of address read from flash memory.

In writing data to a given virtual address, the computer generated address is also divided by the page size to obtain an index to the secondary virtual map in flash memory. The secondary virtual map maps to the 25 primary virtual map, where the primary virtual map block is read; and this is used to map to the physical block that has been addressed where it is read. As this block cannot be rewritten, an unwritten block is identified and written into in the manner previously described, with the original data block marked as deleted. To update the virtual map residing flash 30 memory, essentially the same procedure is followed. The virtual map

block, in its modified form to reflect the new physical location of the address data, is written to an unwritten block in the flash memory and the old block is marked as deleted. The secondary virtual memory in random access memory is changed, as needed, to reflect the change in the primary
5 virtual memory block locations.

Figure 9 is a flow diagram of this operation. The first step in this process is to convert a virtual address to a page number, block 70, and to use the page number to locate in RAM 16 the address, in flash memory 12, of the relevant page block of the virtual map stored in the flash memory,
10 block 21. The page block of the virtual map at this address is read from the flash memory (block 72) and used in the manner previously described to a locate physical address corresponding to the virtual address for a data read or data write operation. In a data write operation, the virtual map page block must be updated, block 73, and the updated page block virtual
15 map is written to a free flash memory physical address location, block 74.

The original flash memory address at which the page block virtual map was located is marked as deleted, block 75, and the RAM memory 16 is updated to point to the virtual to physical map address for the updated map, block 76.

20 The virtual map can be readily reconstructed upon system startup. The virtual maps residing in flash memory are non-volatile and do not require reconstruction. The secondary virtual map residing in volatile random access memory can be reconstructed by scanning, at startup, the block usage map that resides at the top of each unit. Blocks marked as
25 mapped to a virtual address are identified, and the secondary virtual map is constructed accordingly.

While the invention has been described in terms of a single preferred embodiment, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of
30 the appended claims.

WHAT IS CLAIMED IS:

1. A memory management method for a memory in which data can be written only in unwritten physical memory locations and in which a zone of contiguous memory locations can be simultaneously erased, comprising the steps of:

organizing the memory into a plurality of units with each of the units encompassing at least one zone;

organizing each unit into a plurality of blocks, each of said blocks comprised of a plurality of contiguous physical memory locations;

establishing an allocation table for each unit that indicates the status of each block in a unit as active written, unwritten or deleted;

establishing a table to map virtual addresses to physical addresses to physical addresses within a unit;

in writing data to said memory at a virtual address:

(a) mapping said virtual address to a physical block address in a unit;

(b) examining said allocation table for said unit to which said virtual address has been mapped in subparagraph (a) to determine the status of a block at said physical block address as active written or unwritten;

(c) if said block at said physical block address is in an active written status:

(1) examining said allocation table for at least one of said units to identify an unwritten block address;

(2) writing said data into said memory to said unwritten block address;

(3) changing said allocation table for said unit to which said virtual address has been mapped to indicated as deleted said physical block address;

- (4) changing said allocation table for a unit in which said data has been written in paragraph (c)(2) to indicate as active written said unwritten block address where said data has been written;
- (5) changing said table to map virtual addresses to physical addresses within a unit so that said table maps said virtual address to the physical address of said unwritten block in which said data has been written in step (c)(2).

2. A memory management method as in claim 1, including the further steps of:

establishing a transfer unit in said memory in which all blocks of said unit are unwritten;

periodically identifying a selected unit, other than said transfer unit, to be erased;

reading each active written block in said selected unit;

writing each active written block in said selected unit into said transfer unit;

updating said transfer unit allocation table to indicate as written the status of blocks that have been written in the just previous writing step;

erasing said selected block;

updating said table of virtual addresses to physical addresses to denote said selected unit as a transfer unit and said transfer unit with the unit identifier of said selected unit.

3. A memory management method for a memory in which data can be written only in unwritten physical memory locations and in which a zone of contiguous memory locations can be simultaneously erased, comprising the steps of:

storing in said memory a first table that maps virtual addresses to physical addresses;

organizing said first table stored in said memory in segments of page addressable blocks;

storing in a random access memory a second table that maps page addresses to physical addresses of said page addressable blocks in said memory;

changing a page addressable block in said first table stored in said memory by writing a changed page addressable block in an unwritten physical block location; and

updating said second table stored in said random access memory so that it maps the page address of the changed page addressable block to the unwritten physical block location in which said changed page addressable block has been written.

4. A memory management method for a memory in which data can be written only in unwritten physical memory locations and in which a zone of contiguous memory locations can be simultaneously erased, comprising the steps of:

organizing the memory into a plurality of units with each of the units encompassing at least one zone;

organizing each unit into a plurality of blocks, each of said blocks comprised of a plurality of contiguous physical memory locations;

establishing a first table to map virtual addresses to physical addresses within a unit;

storing in said memory said first table organized in segments of page addressable blocks;

storing in a random access memory a second table that maps pages to physical address of said page addressable blocks stored in said memory;

in writing data to said memory at a virtual address:

- (a) deriving a page address from said virtual address;
- (b) mapping said page address to page addressable block in said memory;

- (c) reading a segment of said first table that maps virtual address to physical address at said page addressable block in said memory;
- (d) mapping said virtual address to a physical address;
- (e) if said block at said physical address is in an active written status:
 - (1) writing said data into said memory to an unwritten block address;
 - (2) changing said first table segment so that said first table maps said virtual address to the physical address of the unwritten block in which said data has been written in step (e)(1);
 - (3) writing the changed first table segment from step (e)(2) in an unwritten physical block location in said memory;
 - (4) updating said second table stored in said random access memory so that maps the page address of the changed first table segment of said unwritten physical block location.

5. A memory management method for a memory in which data can be written only in unwritten physical memory locations and in which a zone of contiguous memory locations can be simultaneously erased, comprising the steps of:

organizing the memory into a plurality of units with each of the units encompassing at least one zone;

organizing each unit into a plurality of blocks, each of said blocks comprised of a plurality of contiguous physical memory locations;

establishing an allocation table for each unit that indicates the status of each block in a unit as active written, unwritten or deleted;

establishing a table to map virtual addresses to physical addresses within a unit;

in writing data to said memory at a virtual address:

- (a) mapping said virtual address to a physical block address in a unit;
- (b) examining said allocation table for said unit to which said virtual address has been mapped in subparagraph
 - (a) to determine the status of a block at said physical block address as active written or unwritten;
- (c) if said block at said physical block address is in an active written status:
 - (1) examining said allocation table for at least one of said units to identify an unwritten block address;
 - (2) writing said data into said memory to said unwritten block address;
 - (3) changing said allocation table for said unit to which said virtual address has been mapped to indicate as deleted said physical block address;
 - (4) changing said allocation table for a unit in which said data has been written in paragraph (c)(2) to indicate as active written said unwritten block address where said data has been written;
 - (5) changing said table to map virtual addresses to physical addresses within a unit so that said table maps said virtual address to the physical address of said

unwritten block in which said data has
been written in step (c)(2);

in reading data to said memory at a virtual address:

(d) mapping said virtual address to a physical block address in
a unit;

(e) reading said data from said memory at said physical address.

6. A memory management method as in claim 5, including the further
steps of:

establishing a transfer unit in said memory in which all blocks of
said unit are unwritten;

periodically identifying a selected unit, other than said transfer unit,
to be erased;

reading each active written block in said selected unit;

writing each active written block in said selected unit into said
transfer unit;

updating said transfer unit allocation table to indicate as written the
status of blocks that have been written in the just previous writing step;

erasing said selected block;

updating said table of virtual addresses to physical addresses to
denote said selected unit as a transfer unit and said transfer unit with the
unit identifier of said selected unit.

7. A memory management method for a memory in which data can be
written only in unwritten physical memory locations and in which a zone
of contiguous memory locations can be simultaneously erased, comprising
the steps of:

organizing the memory into a plurality of units with each of the
units encompassing at least one zone;

organizing each unit into a plurality of blocks, each of said blocks
comprised of a plurality of contiguous physical memory locations;

establishing a first table to map virtual addresses to physical addresses within a unit;

storing in said memory said first table organized in segments of page addressable blocks;

storing in a random access memory a second table that maps pages to physical address of said page addressable blocks stored in said memory;

in writing data to said memory at a virtual address:

- (a) deriving a page address from said virtual address;
- (b) mapping said page address to page addressable block in said memory;
- (c) reading a segment of said first table that maps virtual address to physical address at said page addressable block in said memory;
- (d) mapping said virtual address to a physical address;
- (e) if said block at said physical address is in an active written status:
 - (1) writing said data into said memory to an unwritten block address;
 - (2) changing said first table segment so that said first table maps said virtual address to the physical address of the unwritten block in which said data has been written in step (e)(1);
 - (3) writing the changed first table segment from step (e)(2) in an unwritten physical block location in said memory;
 - (4) updating said second table stored in said random access memory so that maps the page address of the changed first table

segment of said unwritten physical block
location;

in reading data to said memory at a virtual address:

- (a) deriving a page address from said virtual address;
- (f) mapping said page address to page addressable block in said memory;
- (g) reading a segment of said first table that maps virtual address to physical address at said page addressable block in said memory;
- (h) mapping said virtual address to a physical address;
- (i) reading said data from said memory at said physical address.

8. A memory management method as in claim 7, including the further steps of:

establishing a transfer unit in said memory in which all blocks of said unit are unwritten;

periodically identifying a selected unit, other than said transfer unit, to be erased;

reading each active written block in said selected unit;

writing each active written block in said selected unit into said transfer unit;

updating said transfer unit allocation table to indicate as written the status of blocks that have been written in the just previous writing step;

erasing said selected block;

updating said table of virtual addresses to physical addresses to denote said selected unit as a transfer unit and said transfer unit with the unit identifier of said selected unit.

FIG.1

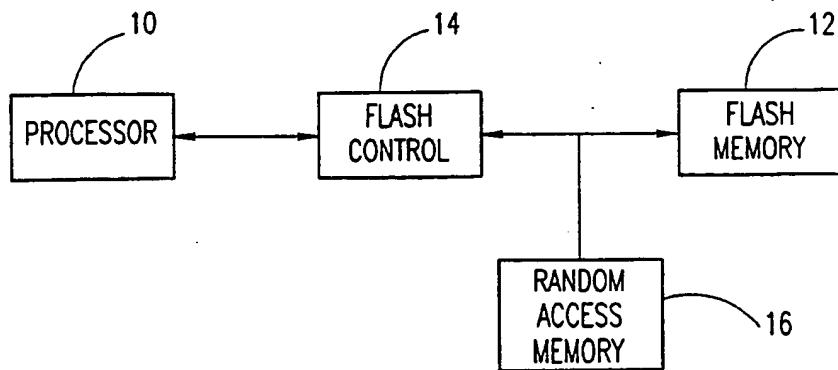


FIG.2

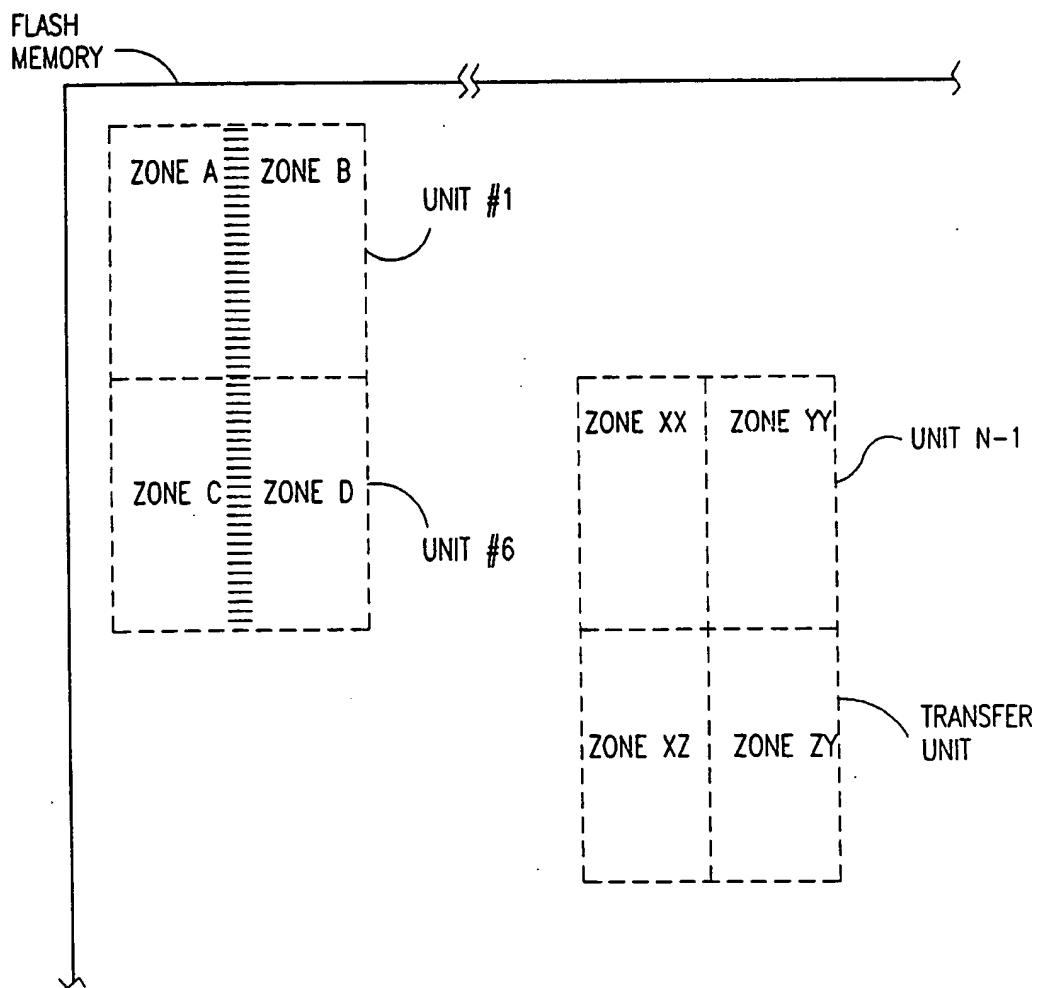


FIG.3

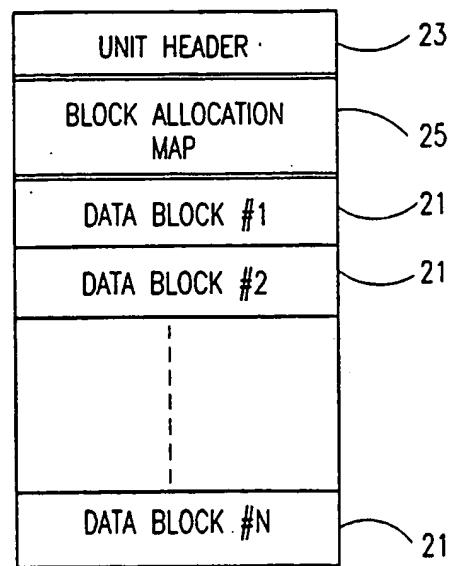


FIG.7

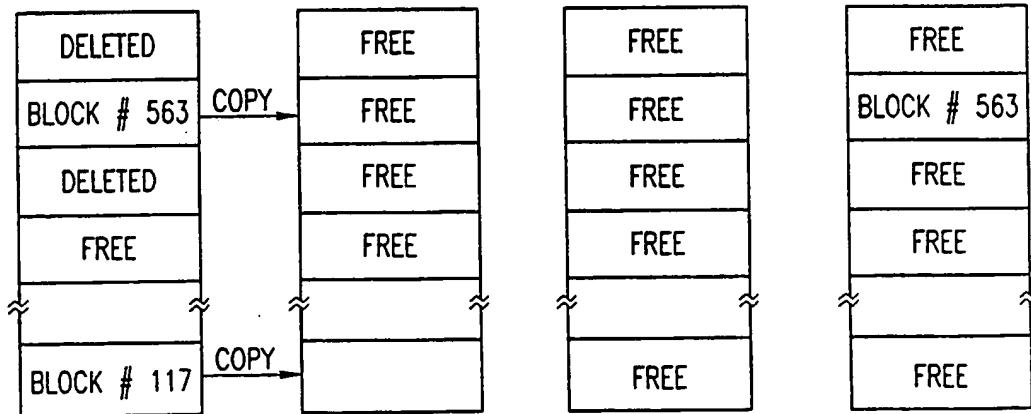
BEFORE

AFTER

LOGICAL UNIT
M

TRANSFER UNIT

TRANSFER UNIT

LOGICAL UNIT
M

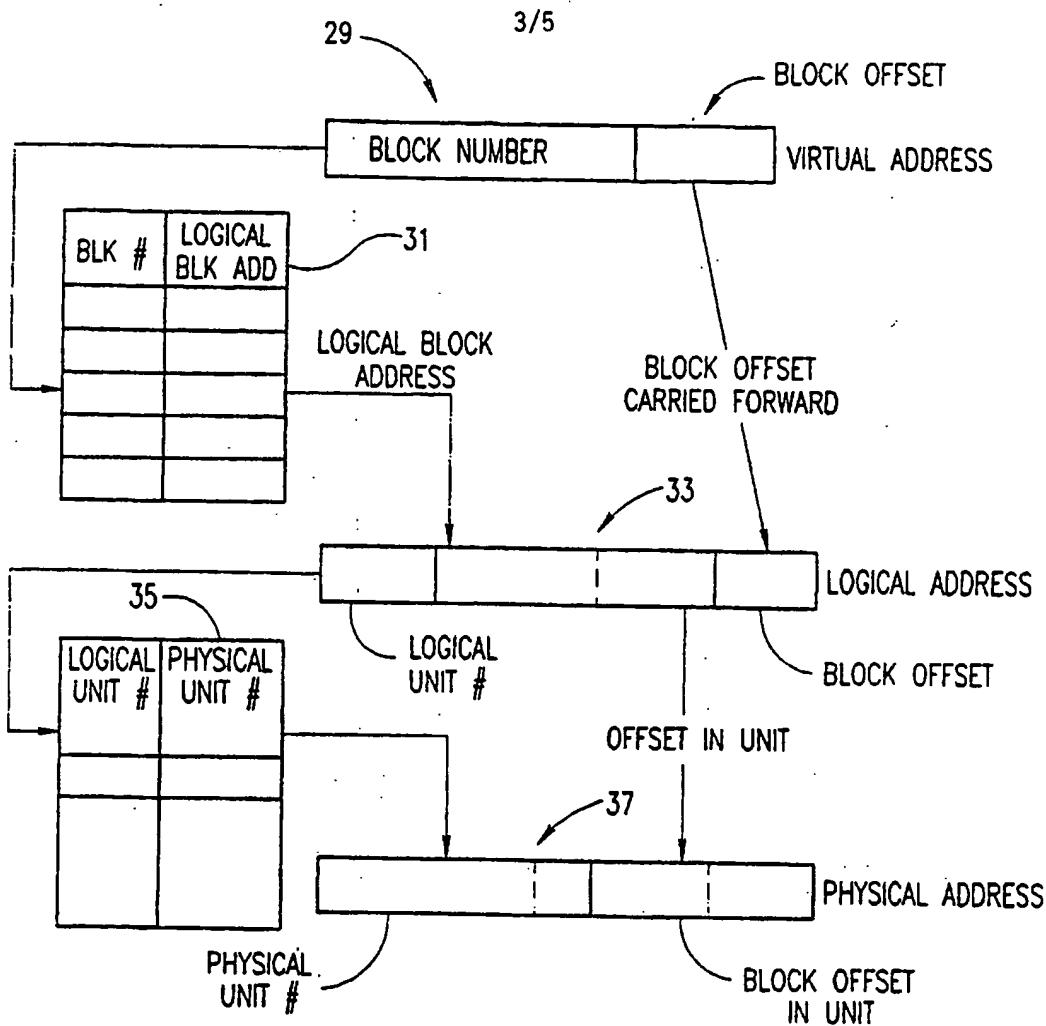


FIG.4

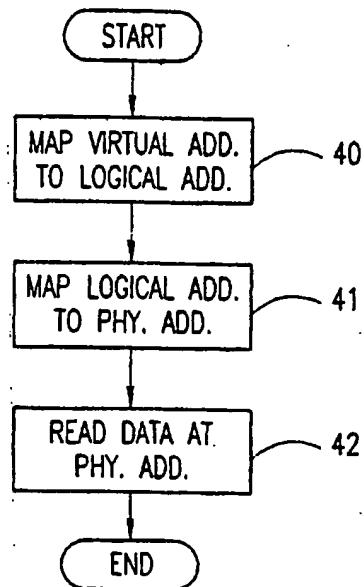
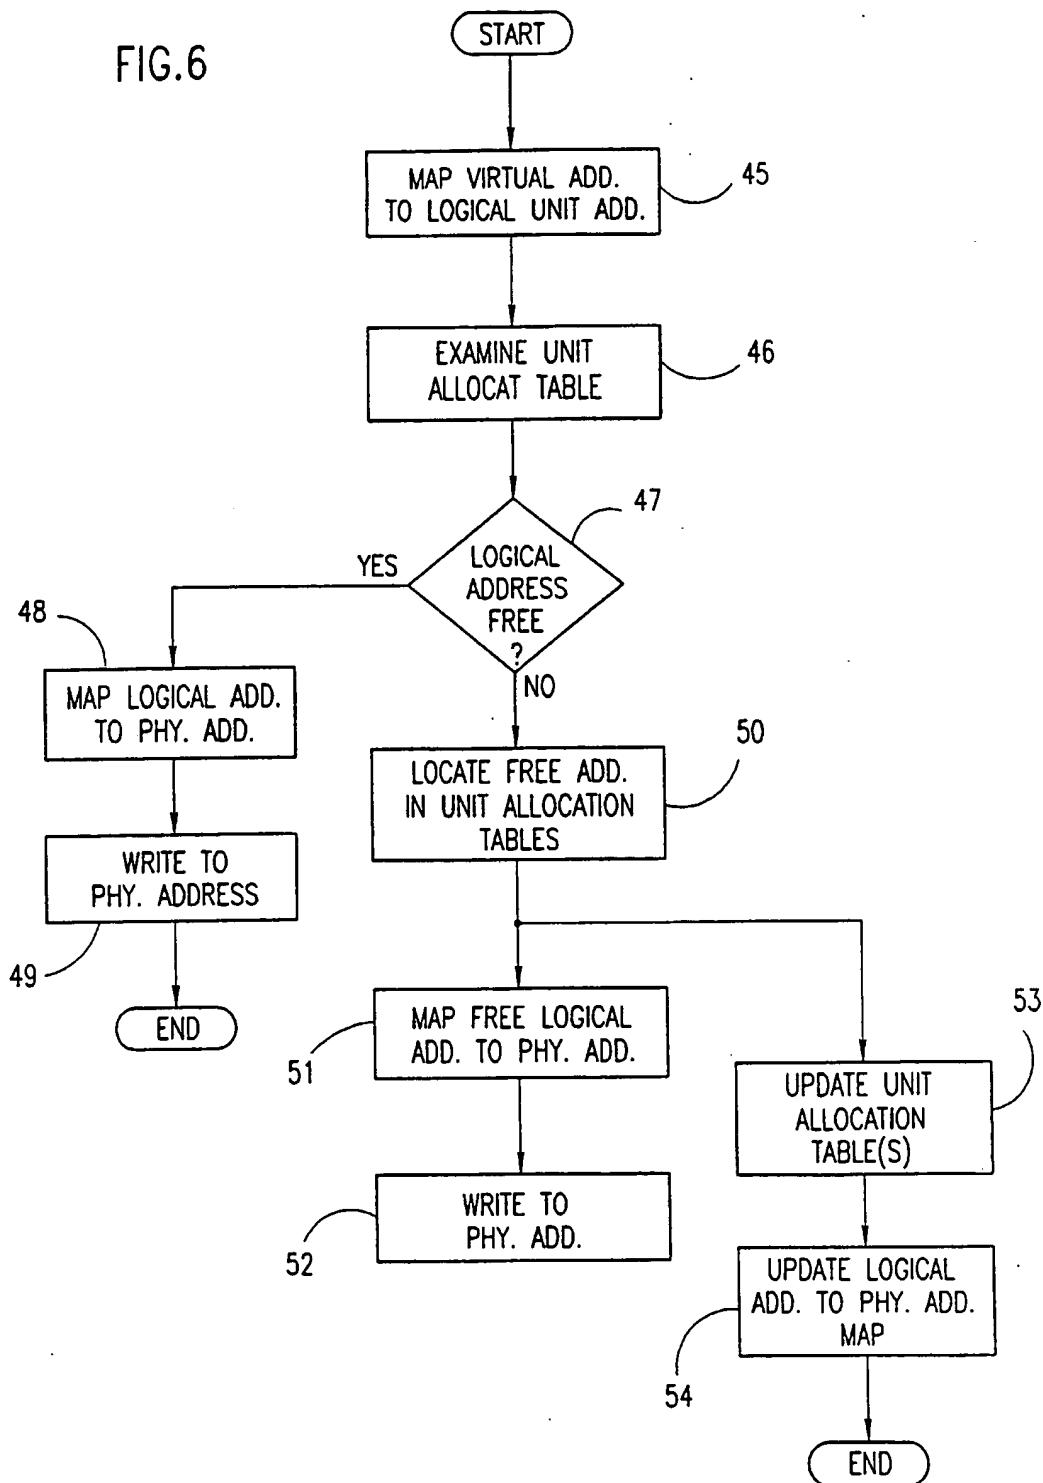
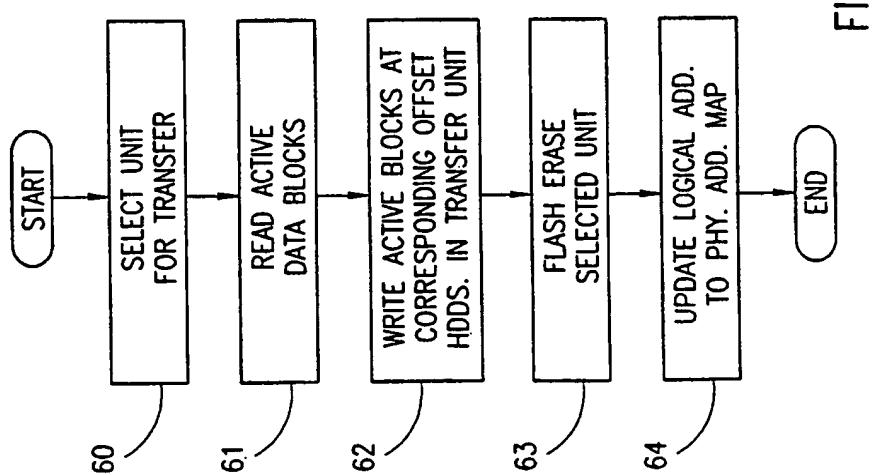
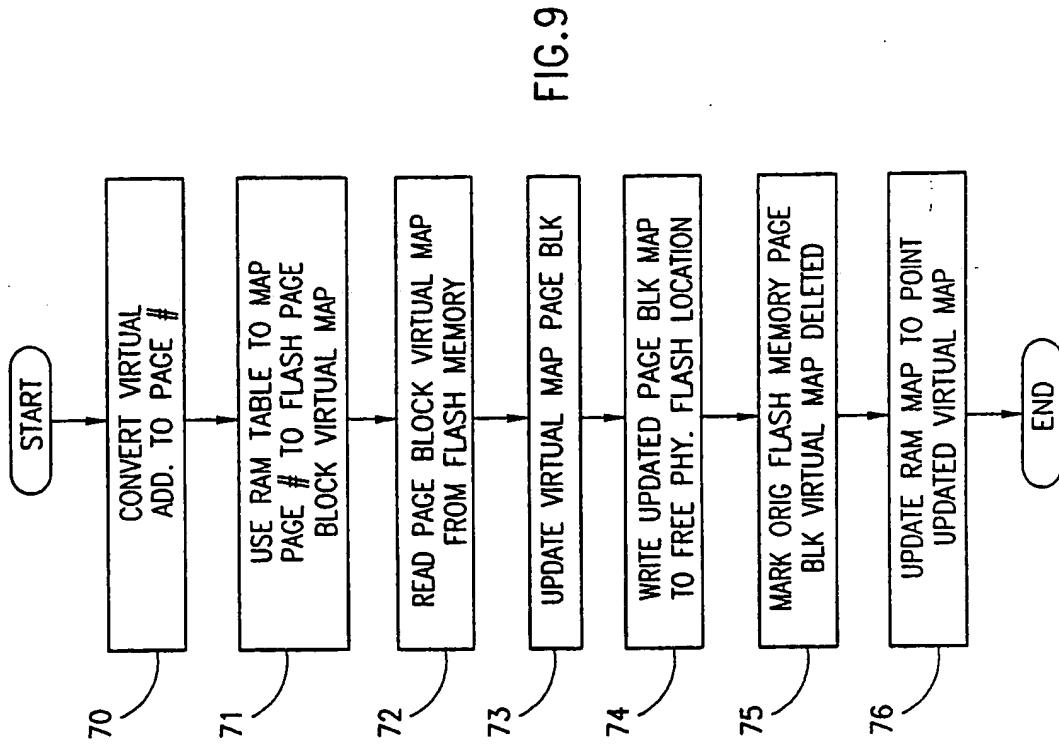


FIG.5

FIG.6





BEST AVAILABLE COPY

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US94/01848

A. CLASSIFICATION OF SUBJECT MATTER

IPC(5) :G06F 12/02

US CL :395/400, 425

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/400, 425 ; 364/200, 900

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 4,511,964 (Goerg et al.) 16 April 1985 See figure 2, column 6, lines 35-53; figure 4, column 8, lines 30-53	1-8
A,P	US, A, 5193,184 (Belsan et al.) 09 March 1993 See the entire document.	1-8
A,P	US, A, 5,210,866 (Milligan et al.) 11 May 1993 See the entire document.	1-8
Y,E	US, A, 5,301,288 (Newman et al.) 05 April 1994 See figure 1, columns 2-3.	1-8

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be part of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Z"	document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

08 JUNE 1994

Date of mailing of the international search report

JUN 24 1994

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Faxsimile No. Not Applicable

Authorized officer
Joseph L. Dixon
B. Handel Jr.

Telephone No. (703) 305-9600